

Structuring Content in the *Façade* Interactive Drama Architecture

Michael Mateas^{1,3} and Andrew Stern^{2,3}

(co-authors listed alphabetically)

¹ College of Computing & Literature, Communication and Culture, Georgia Institute of Technology, michaelm@cc.gatech.edu

² InteractiveStory.net; ³ grandtextauto.org, andrew@interactivestory.org

Abstract

The process of building *Façade*, a first-person, real-time, one-act interactive drama, has involved three major research efforts: designing ways to deconstruct a dramatic narrative into a hierarchy of story and behavior pieces; engineering an AI system that responds to and integrates the player's moment-by-moment interactions to reconstruct a real-time dramatic performance from those pieces; and understanding how to write an engaging, compelling story within this new organizational framework. This paper provides an overview of the process of bringing our interactive drama to life as a coherent, engaging, high agency experience, including the design and programming of thousands of joint dialog behaviors in the reactive planning language ABL, and their higher level organization into a collection of story beats sequenced by a drama manager. The process of iteratively developing the architecture, its languages, authorial idioms, and varieties of story content structures are described. These content structures are designed to intermix to offer players a high degree of responsiveness and narrative agency. We conclude with design and implementation lessons learned and future directions for creating more generative architectures.

Approaching Interactive Story

Stories have rich global, temporal structures whose features can vary both in form and pleasure for audiences. Some stories feature tightly-plotted causal chains of events that may, for example, offer audiences the intrigue of a intricate, unfolding mystery, or the spectacle of an epic historical conflict. By contrast, some stories have sparse, even amorphous event structures, that can offer, for example, the quieter pleasure of following the subtle progression of emotion between two people. The histories of literature, theater, cinema and television demonstrate that many types of story structures can be pleasurable for audiences; the challenge for researchers and artists is determining how traditional story forms can be adapted for interactivity.

Interactive experiences have several identifiable features of their own, such as immersion, agency, and transformation (Murray 1997), each offering particular pleasures for interactors, and varying compatibility with

story. For many artists and researchers, *agency* is often considered the holy grail of interactive story pleasures, perhaps because it offers players the most substantial influence on the overall structure of the experience. Agency is also the most challenging to implement, exactly because it requires the system to dynamically assemble a story structure that incorporates the unpredictable actions of the player. This suggests that stories with *looser, sparser* event structures (plots) will be easier to implement in an interactive medium (require less generativity), and thus should be a good starting point for interactive story researchers and artists. Note that such stories can be as pleasurable as tightly-plotted ones, just in different ways.

When designing an interactive story architecture and its content structures, the design choices made will influence the types of stories that can be built within it, and greatly affect the likelihood of ultimately creating pleasurable experiences for players. With this in mind, the *Façade* architecture was designed with features intended for building experiences with *high agency*, and with enough narrative intelligence (NI) (Mateas and Sengers 2002) to construct character-rich, causally-sparse yet coherent plots. Further, we chose to implement several NI features particular to *theatrical drama*, a powerful story form historically shown to be compatible with sparse plots, compensated for by rich *emotional expression* from its characters. Additional features therefore important to support include characters with a strong sense of *immediacy* and *presence*, whose very *aliveness* results in the audience experiencing a sensation of danger or unpredictability, that anything is possible.

This paper presents *Façade's* solution to the tension inherent between game and story, some organizing principles allowing us to move away from traditional branching narrative structures, and an overview of *Façade's* architecture, combined with how its content is structured and applied. We describe *Façade's* atomic unit of dramatic performance, the joint dialog behavior, the variety of its applications within the drama, their organization into story beats that afford sparse but coherent plots, and their integration with sets of forward-chaining natural language processing (NLP) rules offering players a high degree of emotional expression. We conclude with design and implementation lessons learned, and future directions for creating more generative architectures. All of this discussion is guided by our primary design goal: to create an architecture for, and

working example of, high agency, emotionally expressive interactive drama.

Resolving Game Versus Story

Today's most pleasurable high agency interactive experiences are games, because the mechanics of game agency are well understood and reasonably straightforward to implement. Player moves such as running, jumping or shooting, playing a card, or moving a pawn directly cause scores, stats, levels or abstract game-piece configurations to change. (Simulations of physical environments and resource-bound systems have more complex state, but can still be represented numerically in understood ways.) However, to date, a high agency interactive story has yet to be built. Existing game design and technology approaches, that focus on the feedback loop between player interaction and relatively simple numeric state, seem inappropriate for modeling the player's effect on *story structure*, whose complex global constraints seem much richer than can be captured by a set of numeric counters or game pieces.

Our solution to this long-time conundrum is to recast interactions within a story world in terms of abstract *social games*. At a high level, these games are organized around a numeric "score", such as the affinity between a character and the player. However, unlike traditional games in which there is a fairly direct connection between player interaction (e.g. pushing a button to fire a gun) and score state (e.g. a decrease in the health of a monster), in our social games several levels of abstraction may separate atomic player interactions from changes in social "score". Instead of jumping over obstacles or firing a gun, in *Façade* players fire off a variety of *discourse acts*, in natural language, such as praise, criticism, flirtation and provocation. While these discourse acts will generate immediate reactions from the characters, it may take story-context-specific patterns of discourse acts to influence the social game score. Further, the score is not directly communicated to the player via numbers or sliders, but rather via enriched, theatrically dramatic performance.

As a friend invited over for drinks at a make-or-break moment in the collapsing marriage of the protagonists Grace and Trip, the player unwittingly becomes an antagonist of sorts, forced by Grace and Trip into playing psychological "head games" with them (Berne 1964). During the first part of the story, Grace and Trip interpret all of the player's discourse acts in terms of a zero-sum *affinity game* that determines whose side Trip and Grace currently believe the player to be on. Simultaneously, the *hot-button game* is occurring, in which the player can trigger incendiary topics such as sex or divorce, progressing through tiers to gain more character and backstory information, and if pushed too far on a topic, affinity reversals. The second part of the story is organized around the *therapy game*, where the player is (purposefully or not) potentially increasing each characters' degree of self-realization about their own

problems, represented internally as a series of counters. Additionally, the system keeps track of the overall story *tension level*, which is affected by player moves in the various social games. Every change in each game's state is performed by Grace and Trip in emotionally expressive, dramatic ways. On the whole, because their attitudes, levels of self-awareness, and overall tension are regularly progressing, the experience takes on the form and aesthetic of a loosely-plotted domestic drama.



Figure 1. Grace and Trip in *Façade*, viewed from the player's first-person perspective.

Richness Through Coherent Intermixing

Even with a design solution in hand for resolving the tension between game and story, an organizing principle is required to break away from the constraints of traditional branching narrative structures, to avoid the combinatorial explosion that occurs with complex causal event chains (Crawford 1989). Our approach to this in *Façade* is twofold: first, we divide the narrative into *multiple fronts of progression*, often causally independent, only occasionally interdependent. Second, we build a variety of *narrative sequencers* to sequence these multiple narrative progressions. These sequencers operate in parallel and can coherently intermix their performances with one another.

Façade's architecture and content structure are two sides of the same coin, and will be described in tandem; along the way we will describe how the coherent intermixing is achieved.

Architecture and Content Structure

The *Façade* architecture consists of characters written in the reactive-planning language ABL, a drama manager that sequences dramatic beats, a forward-chaining rule system for understanding and interpreting natural language and gestural input from the player, and an animation engine that performs real-time non-photorealistic rendering, spoken dialog, music and sound, driven by and providing sensing data to the ABL behaviors (Mateas & Stern 2004a; Mateas & Stern 2004b;

Mateas & Stern 2003a; Mateas & Stern 2003b; Mateas & Stern 2000).

The narrative sequencers for the social games are written in ABL, often taking advantage of ABL's support for reflection in the form of meta-behaviors that can modify the runtime state of other behaviors. The highest level narrative sequencer, the drama manager, sequences dramatic beats which are described in a custom drama management language.

Beats, Beat Goals and Beat Mix-ins

Façade's primary narrative sequencing occurs within a *beat*, inspired by the smallest unit of dramatic action in the theory of dramatic writing (McKee 1997); however *Façade* beats ended up being larger structures than the canonical beats of dramatic writing. A *Façade* beat is comprised of anywhere from 10 to 100 *joint dialog behaviors* (jdb's), written in ABL. Each beat is in turn a narrative sequencer, responsible for sequencing a subset of its jdb's in response to player interaction. Only one beat is active at any time. A jdb, *Façade*'s atomic unit of dramatic action (and closer to the canonical beat of dramatic writing) consists of a tightly coordinated, dramatic exchange of 1 to 5 lines of dialog between Grace and Trip, typically lasting a few seconds. Jdbs consist of 40 to 200 lines of ABL code. A beat's jdbs are organized around a common narrative goal, such as a brief conflict about a topic, like Grace's obsession with redecorating, or the revelation of an important secret, like Trip's attempt to force Grace to enjoy their second honeymoon in Italy. Each jdb is capable of changing one or more values of story state, such as the affinity game's spectrum value, or any of the therapy game's self-revelation progression counters, or the overall story tension level. In the first part of the story, the within-beat narrative sequencer implements the affinity game; the topic of the beat is organized as an instance of the affinity game.

There are two typical uses of jdbs within beats: as *beat goals* and *beat mix-ins*. A beat consists of a canonical sequence of narrative goals called beat goals. The typical canonical sequence consists of a transition-in goal that provides a narrative transition into the beat (e.g. bringing up a new topic, perhaps connecting it to the previous topic), several body goals that accomplish the beat (in affinity game beats, the body goals establish topic-specific conflicts between Grace and Trip that force the player to choose sides), a wait goal in which Grace and Trip wait for the player to respond to the head game established by the beat, and a default transition-out that transitions out of the beat in the event of no player interaction. In general, transition-out goals both reveal information and communicate how the player's action within the beat has changed the affinity dynamic.

The canonical beat goal sequence captures how the beat would play out in the absence of interaction. In addition to the beat goals, there are a set of handler meta-behaviors that wait for specific NLP interpretations of player discourse acts, and modify the canonical sequence in

response, typically using beat mix-ins. That is, the handler logic implements the custom narrative sequencer for the beat. Beat mix-in jdbs are beat-specific reactions used to respond to player actions and connect the interaction back to the canonical sequence. Handlers are responsible both for potentially adding, removing and re-ordering future beat goals, as well as interjecting beat mix-ins into the canonical sequence. By factoring the narrative sequencing logic and the beat goals in this way, we avoid having to manually unwind the sequencing logic into the beat goal jdbs themselves.

For *Façade*, an experience that lasts ~20 minutes and requires several replays to see all of the content available (any one runthrough performs at most 25% of the total content available), we authored ~2500 jdbs. Approximately 66% of those 2500 are in beat goals and beat mix-ins, organized into 27 distinct beats, of which ~15 are encountered by the player in any one runthrough (see the drama management section further below).

Global Mix-in Progressions

Another type of narrative sequencer, that operates in parallel to and can intermix with beat goals and beat mix-ins, are *global mix-ins*. (How coherent intermixing is achieved is described in a later section.) Each category of global mix-in has three tiers, progressively digging deeper into a topic; advancement of tiers is caused by player interaction, such as referring to the topic. Each tier in the progression is constructed from one or more jdbs, just like beat goals or beat mix-ins. They are focused on satellite topics such as marriage, divorce, sex, therapy, or about objects such as the furniture, drinks, their wedding photo, the brass bull, or the view, or as generic reactions to praise, criticism, flirtations, oppositions and the like. Additionally, there are a variety of generic deflection and recovery global mix-ins for responding to overly confusing or inappropriate input from the player. In total there are ~20 instances of this type of narrative sequencer in *Façade*, comprising about 33% of the total ~2500 jdbs.

Drama Management (Beat Sequencing)

The coarsest narrative sequencing in *Façade* occurs in the drama manager, or *beat sequencer*. This lies dormant

PlayerArrives, TripGreetsPlayer, PlayerEntersTripGetsGrace, GraceGreetsPlayer, ArgueOverRedecorating, ExplainDatingAnniversary, ArgueOverItalyVacation, FixDrinksArgument, PhoneCallFromParents, TransitionToTension2, GraceStormsToKitchen, PlayerFollowsGraceToKitchen, GraceReturnsFromKitchen, TripStormsToKitchen, PlayerFollowsTripToKitchen, TripReturnsFromKitchen, TripReenactsProposal, BigBlowupCrisis, PostCrisis, TherapyGame, Revelations-Buildup, Revelations, EndingNoRevelations, Ending-SelfRevelationsOnly, EndingRelationshipRevelationsOnly, EndingBothNotFullySelfAware, EndingBothSelfAware

Table 1. The names of *Façade*'s 27 beats.

most of the time, only active when the current beat is finished or is aborted (by the beat's own decision, or by a global mix-in). It is at the beat sequencing level where causal dependence between major events is handled – that is, where high-level plot decisions are made.

In a beat sequencing language the author annotates each beat with selection knowledge consisting of preconditions, weights, weight tests, priorities, priority tests, and story value effects – the overall tension level, in *Façade*'s case. Given a collection of beats represented in the beat language, such as the 27 listed in Table 1, the beat sequencer selects the next beat to be performed. The unused beat whose preconditions are satisfied and whose story tension effects most closely match the near-term trajectory of an author-specified story tension arc (in *Façade*, an Aristotelian tension arc) is the one chosen; weights and priorities also influence the decision. (Mateas & Stern 2003a)

Subsequent sections on Context Intermixing and Failures and Successes further discuss beat sequencing.

Long-term Autonomous Mix-in Behaviors

Long-term autonomous behaviors, such as fixing drinks and sipping them over time, or carrying around and compulsively playing with an advice ball toy, last longer than a 60-second beat or a 10-second global mix-in. While perhaps performing only a minor narrative function, occasionally mixing in a jdb into the current beat (comprising only 1% of *Façade*'s jdbs), they contribute a great deal to the appearance of intelligence in the characters, by having them perform extended, coherent series of low-level actions in the background over the course of many minutes, across several beat boundaries. By simultaneously performing completely autonomous behaviors and joint behaviors, *Façade* characters are a hybrid between the “one-mind” and “many-mind” extremes of approaches to agent coordination, becoming in effect “multi-mind” agents (Mateas & Stern 2004a).

Strategies for Coherent Intermixing

Since global mix-ins for the hot-button game are sequenced among beat goals/mix-ins for the affinity game, which both operate in parallel with the drama manager that is occasionally progressing overall story tension, several strategies are needed to maintain coherency, both in terms of discourse management and narrative flow.

First, global mix-in progressions are written to be causally independent of any beats' narrative flow. For example, while quibbling about their second honeymoon in Italy, or arguing about what type of drinks Trip should serve (affinity game beats, chosen by the drama manager), it is safe to mix in dialog about, for example, sex, or the wedding photo (hot-button game mix-ins, triggered by a player's reference to their topics). Each mix-in's dialog is written and voice-acted as if they are slightly tangential

topics that are being jugged into the flow of conversation (“Oh, that photo, yeah, it's really...”).

At the discourse level, mechanisms exist for smoothly handling such interruptions. During a beat goal, such as Trip's reminiscing about the food in Italy, if a global mix-in is triggered, such as the player picking up (referring to) the brass bull (a gift from Trip's lover), the current Italy beat goal will immediately stop mid-performance, and the brass bull global mix-in will begin performing, at whichever tier that hot-button game has already progressed to. At the time of interruption, if the Italy beat goal had not yet passed its *gist point*, which is an author-determined point in a beat goal's jdbs, it will need to be repeated when the global mix-in completes. Short, alternate uninterruptible dialog is authored for each beat goal for that purpose. Also, each beat goal has a *reestablish* jdb that gets performed if returning to the beat from a global mix-in (“So, I was going to say, about Italy...”). Mix-in's themselves can be interrupted by other mix-in's, but if so, are not repeated as beat goals are.

With only a few exceptions, the narratives of affinity game beats themselves are also designed to be causally independent of one another, relating to the “sparse plot” characterization made earlier. For example, it does not matter which order Grace and Trip argue about Italy, their parents, redecorating, fixing drinks, or their dating anniversary. When beat sequencing, this allows the drama manager to prefer sequencing any beats related to past topics brought up by the player. Likewise, hot-button mix-ins can be safely triggered in any order, into almost any beat at any time.

However, great authorial effort was taken to make the *tone* of each beat goal/mix-in and global mix-in match each other during performance. Most jdbs are authored with 3 to 5 alternates for expressing its narrative content at different combinations of player affinity and tension level. These include variations in word choice, voice-acting, emotion, gesture, and appropriate variation of information revealed. By having the tone of hot-button global mix-ins and affinity game beat goals/mix-ins always match each other, players often perceive them as causally related, even though they are not. Additionally, for any one tone, most jdbs are authored with 2 to 4 dialog alternates, equivalent in narrative functionality but helping create a sense of freshness and non-roboticness in the characters between runthroughs of the drama.

Evaluating Agency

The structure of narrative content in *Façade*, described in the previous sections, is intended to afford high agency for players, a primary pleasure of interactive experiences. In this section we identify two types of agency: local and global, and attempt to evaluate the degree of their existence in *Façade*.

Local Agency

When the player's actions cause immediate, context-specific, meaningful reactions from the system, we call this *local agency*. Furthermore, the greater the range of actions the player can take, that is, the more expressive the interface, then the richer the local agency (again, if the responses are meaningful).

Façade offers players a continuous, open-ended natural language interface, as well as physical actions and gestures such as navigation, picking up objects, hugging and kissing. The millions of potential player inputs are mapped, using hundreds of authored *forward-chaining NLU rules*, into one or more of ~30 parameterized discourse acts (DA's) such as praise, exclamation, topic references, and explanations. Another set of rules called *context proposers* then interpret these DA's in context-specific ways, such as agreement, disagreement, alliance, or provocation (Mateas and Stern 2004b).

Ideally there would be immediate, meaningful, context-specific responses available at all times for all DA's. In the actual implementation of *Façade*, in our estimation this ideal is reached ~25% of the time, where the player has a satisfying degree of real-time control over Grace and Trip's emotional state, affinity to the player, which topic is being debated, what information is being revealed, and the current tension level. But more often, ~40% of the time, only a partial ideal is reached: the mapping/interpretation from DA to reaction is coarser, the responses are more generic and/or not as immediate. Furthermore, ~25% of the time even shallower reactivity occurs, and ~10% of the time there is little or no reactivity. These varying levels of local agency are sometimes grouped together in temporal clusters, but also have the potential to shift on a moment-by-moment basis.

There are two main reasons for these varying levels of local agency. First, from a design perspective, at certain points in the overall experience it becomes necessary to funnel the potential directions of the narrative in authorially preferred directions, to ensure dramatic pacing and progress. Second, and more often the case, a lack of local agency is due to limitations in how much narrative content was authored (see the Failures section below).

Global Agency

The player has *global agency* when the global shape of the experience is determined by player action. In *Façade* this would mean that the final ending of the story, and the particulars of the narrative arc that lead to that ending, are determined in a smooth and continuous fashion by what the player does, and that at the end of the experience the player can understand how her actions led to this storyline.

Façade attempts to achieve global agency in a few ways. First, beat sequencing (i.e., high level plot) can be influenced by what topics the player refers to; the sequencing can vary within the number of allowed

permutations of the beats' preconditions and tension-arc-matching requirements. Even with only 27 beats in the system, technically there are thousands of different beat sequences possible; however, since most beats are causally independent, the number of *meaningfully different* beat sequences are few.

More significant than variations of beat sequences ("what" happened) are variations within beats and global mix-in progressions ("how" it happened). A variety of patterns and dynamics are possible within the affinity, hot-button and therapy games over the course of the experience; in fact these patterns are monitored by the system and remarked upon in dramatic recapitulations in the BigBlowupCrisis beat halfway through the drama, and in the RevelationsBuildup beat at the climax of the drama. A calculus of the final "scores" of the various social games is used to determine which of five ending beats gets sequenced, ranging from either Grace or Trip revealing one or more big hidden secrets and then deciding to break up and leave, or both of them too afraid to do anything, or both them realizing so much about themselves and each other that they decide to stay together.

Failures and Successes

During the production of *Façade*, within our "limited" authoring effort (beyond the building of the architecture, *Façade* required ~3 person years of just authoring, which is more than a typical art/research project but far less than a typical game industry project) we made the tradeoff to support a significant degree of local agency, which in the end came at the expense of global agency. Combined with the reality that the time required to design and author jdbs is substantial, only 27 beats were created in the end, resulting in far lower global agency than we initially hoped for. As a result, we feel we did not take full advantage of the power of the drama manager's capabilities.

Furthermore, because the specification of each joint dialog behavior – spoken dialog, staging directions, emotion and gesture performance – requires a great deal of authoring and is not automatically generated by higher-level behaviors or authoring tools, we are limited to the permutations of hand-authored, intermixable content. *Façade* is not generating sentences – although it is generating sequences.

A major challenge we encountered, that we believe *Façade* falls short on, is always clearly communicating the state of the social games to the player. With traditional games, it is straightforward to tell players the game state: display a numeric score, or show the character physically at a higher platform, or display the current arrangement of game pieces. But when the "game" is ostensibly happening inside of the characters' heads, and if we intend to maintain a theatrical, performative aesthetic (and not display internal feelings via stats and slider bars, ala *The Sims*), it becomes a significant challenge. In our estimation *Façade* succeeds better at communicating the

state of the simpler affinity and hot-button games than the more complex therapy game.

Another major challenge was managing the player's expectations, raised by the existence of an open-ended natural language interface. We anticipated natural language understanding failures, which in informal evaluations of *Façade* to date, occur ~30% of the time on average. This tradeoff was intentional, since we wanted to better understand the new pleasures that natural language can offer when it *succeeds*, which in *Façade* we found occurs ~70% of the time, either partially or fully.

In our estimation, a success of *Façade* is the integration of the beat goal/mix-in, global mix-in and drama manager narrative sequencers, with an expressive natural language interface, context-specific natural language processing, and expressive real-time rendered character animation. We feel the overall effect makes some progress towards our original design goals of creating a sense of the immediacy, presence, and aliveness in the characters required for theatrical drama.

Certain aspects of our drama's design help make *Façade* a pleasurable interactive experience, while others hurt. It helps to have *two* tightly-coordinated non-player characters who can believably keep dramatic action happening, in the event that the player stops interacting or acts uncooperatively. In fact, the fast pace of Grace and Trip's dialog performance discourages lengthy natural language inputs from the player. By design, Grace and Trip are *self-absorbed*, allowing them to occasionally believably ignore unrecognized or unhandleable player actions. Creating a loose, sparsely plotted story afforded greater local agency, but provided fewer opportunities for global agency. However, the richness of content variation, and at least moderate degree of global agency achieved, does encourage replay.

The huge domain of the drama, a marriage falling apart, arguably hurt the success of the overall experience, in that it overly raised players' expectations of the characters' intelligence, psychological complexity, and language competence. As expected, the system cannot understand, nor has authored reactions for, many reasonable player utterances. The large domain often requires mapping millions of potential surface texts to just a few discourse acts, which can feel muddy or overly coarse to the player. Also, continuous real-time interaction, versus discrete (turn-taking) and/or non-real-time interaction, added a great deal of additional complexity and authoring burden.

Future Directions

In order to relieve the high authoring burden encountered in building *Façade*, particularly in writing the thousands of joint dialog behaviors required for an interactive drama with significant player agency, an even more *generative* approach is required. One approach is to create a higher-level authoring tool, which compiles a more abstract specification of narrative behavior into a set of

corresponding joint dialog ABL behaviors. Greater authoring productivity would allow for increased local and global agency, since higher numbers of performance behaviors could be created more quickly and efficiently.

Such an authoring tool would substantially increase the possibility for artists with minimal programming skills to help build high agency interactive dramas. Short of that, developers with significant aptitude and experience in programming, dramatic writing and game design will be required, of which there are currently few, and few educational curricula to create them.

Additional approaches include implementing more low-level support for managing and generating emotion-rich, dramatic conversation behaviors, such as integrating the Em emotion bookkeeping system (Reilly 1997), automated non-verbal conversation gestures (Vilhjálmsson 2003), or procedural body animation (Perlin 2004).

References

- Berne, E. 1964. *Games People Play*. New York: Grove Press.
- Crawford, C. 1989. Indirection. *Journal of Computer Game Design, Volume 3*. http://www.erasmatazz.com/library/JCGD_Volume_3/JCGD_Volume_3_Index.html
- Mateas, M. and Sengers, P. (Eds.). 2002. *Narrative Intelligence*. Amsterdam: John Benjamins.
- Mateas, M. and Stern, A. 2004a. A Behavior Language: Joint Action and Behavioral Idioms. In Predinger, H. and Ishioka, M. (Eds.), *Life-like Characters: Tools, Affective Functions and Applications*, Springer 2004.
- Mateas, M. and Stern, A. 2004b. Natural Language Understanding in *Façade*: Surface Text Processing. In *Proceedings of Technologies for Interactive Digital Storytelling and Entertainment*, Darmstadt, Germany, June 2004.
- Mateas, M. and Stern, A. 2003a. Integrating plot, character and natural language processing in the interactive drama *Façade*. In *1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE '03)*, Darmstadt, Germany, March 24 – 26, 2003.
- Mateas, M. and Stern, A. 2003b. *Façade*: an experiment in building a fully-realized interactive drama, *Game Developers Conference (GDC '03)*, San Jose, CA, USA, March 4 – 8, 2003.
- Mateas, M. and Stern, A. 2000. Towards Integrating Plot and Character for Interactive Drama. In *Working notes of the Social Intelligent Agents: The Human in the Loop Symposium*. AAAI Fall Symposium Series. Menlo Park, CA.: AAAI Press.
- McKee, R. 1997. *Story: Substance, Structure, Style, and the Principles of Screenwriting*. New York: HarperCollins.
- Murray, J. 1997. *Hamlet on the Holodeck*. New York: Free Press. 1998, Cambridge, MA: MIT Press.
- Perlin, K. 2004. Emotive Actors. Media Research Laboratory, NYU. <http://mrl.nyu.edu/~perlin/experiments/emotive-actors>
- Reilly, S. 1997. Believable Social and Emotional Agents. Ph.D. thesis, Tech report CMU-CS-96-138, Carnegie Mellon University.
- Vilhjálmsson, H. 2003. Avatar Augmented Online Conversation. Ph.D. thesis, MIT Media Lab.